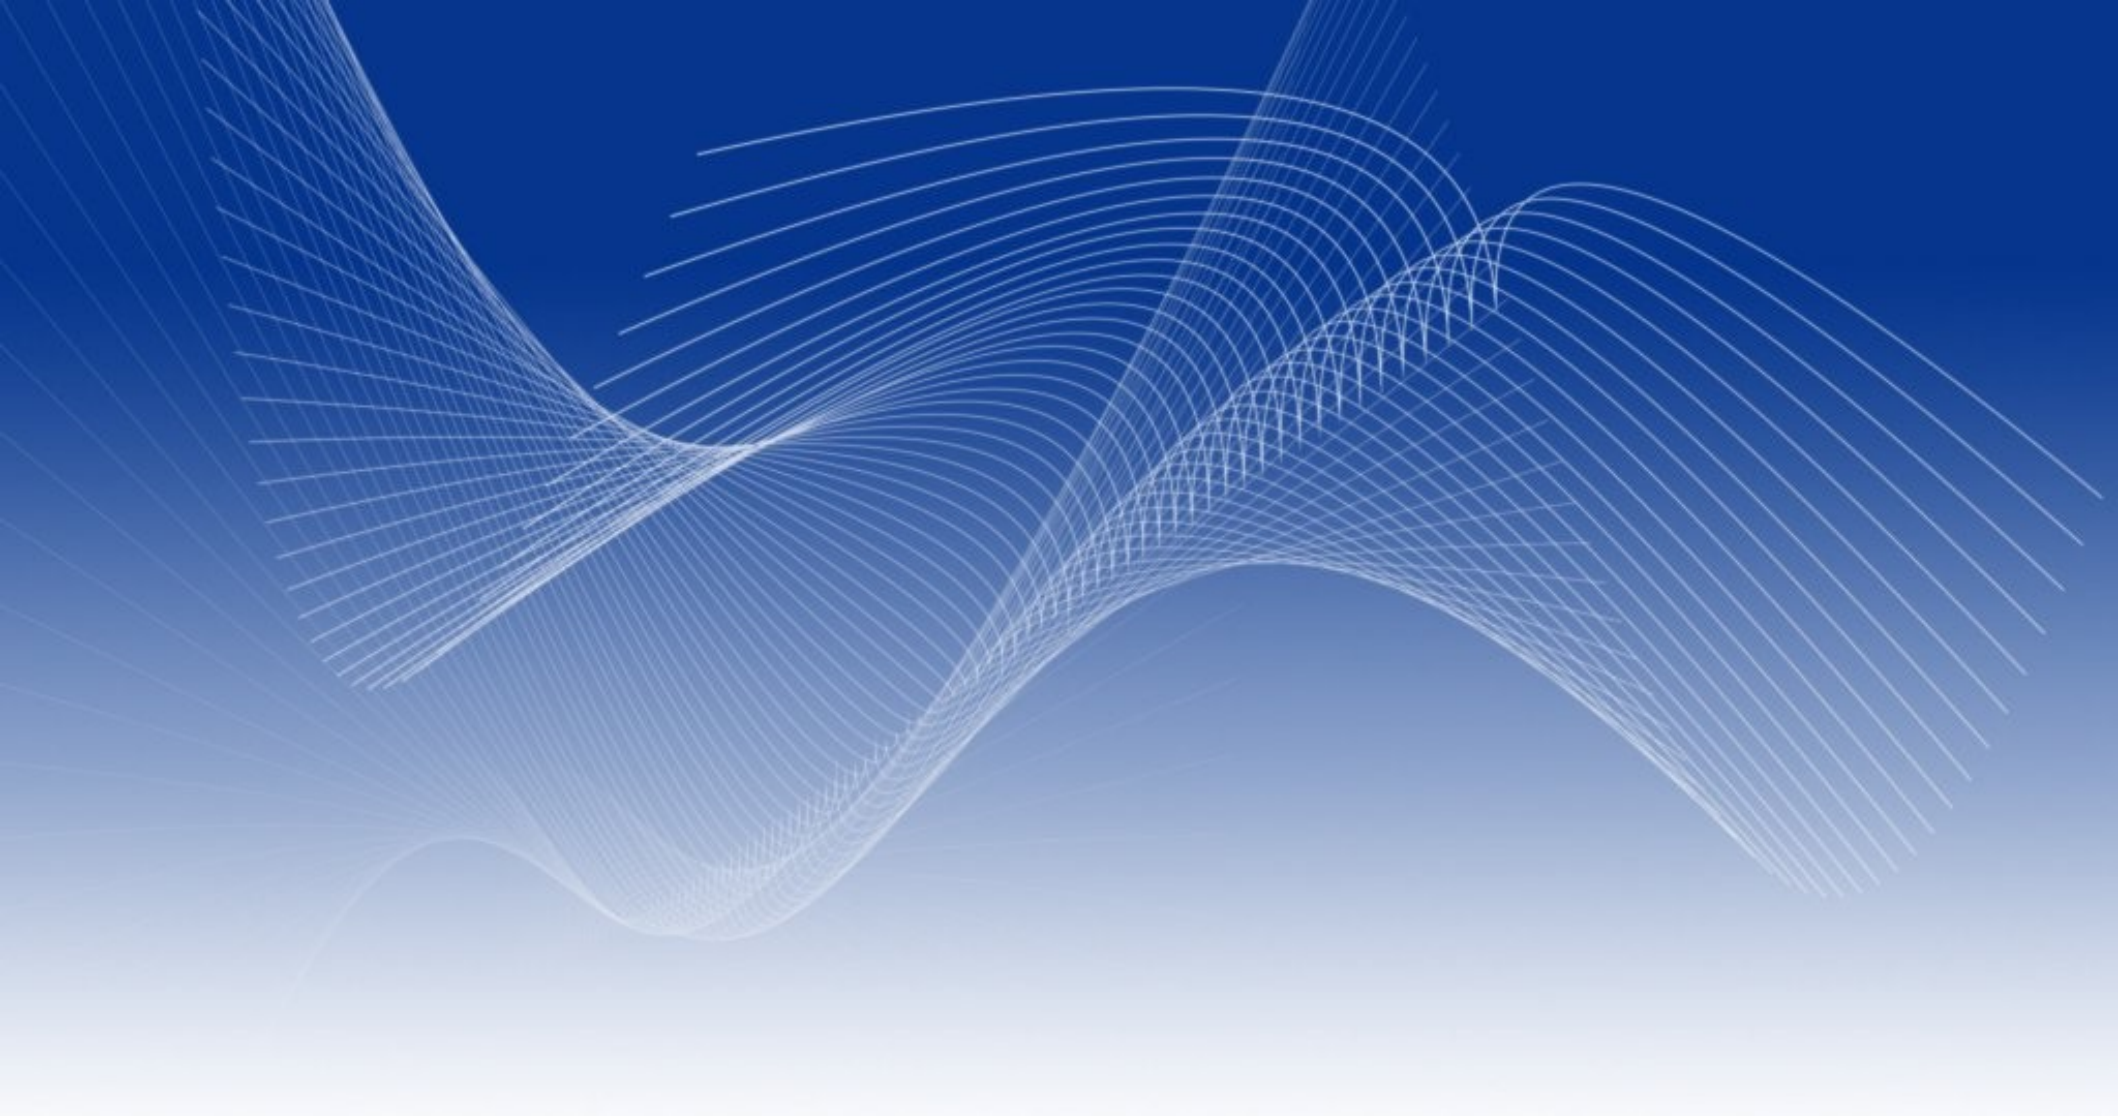

End to End Information Infrastructure & Services in Open Science Grid (OSG)

Anthony Tiradani
Fermilab

Definitions and Scope

- Information Types:
  - Monitoring: Information that gives a view on resource health and current availability. Example: RSV
  - Accounting: Information that gives a view of what resources are being/has been used, for how long, and by whom. Example: Gratia
  - Discovery: Information that provides a view of what services are offered and what those services look like. Example: RESS/BDII
- This discussion will deal with Discovery information. How is it generated and published? Who consumes it?
- Note: The current implementation of the Discovery Information contains data that can cross definitions. There are groups that use it for other purposes (Installed Capacity).

Current Information Flow (Simplified)

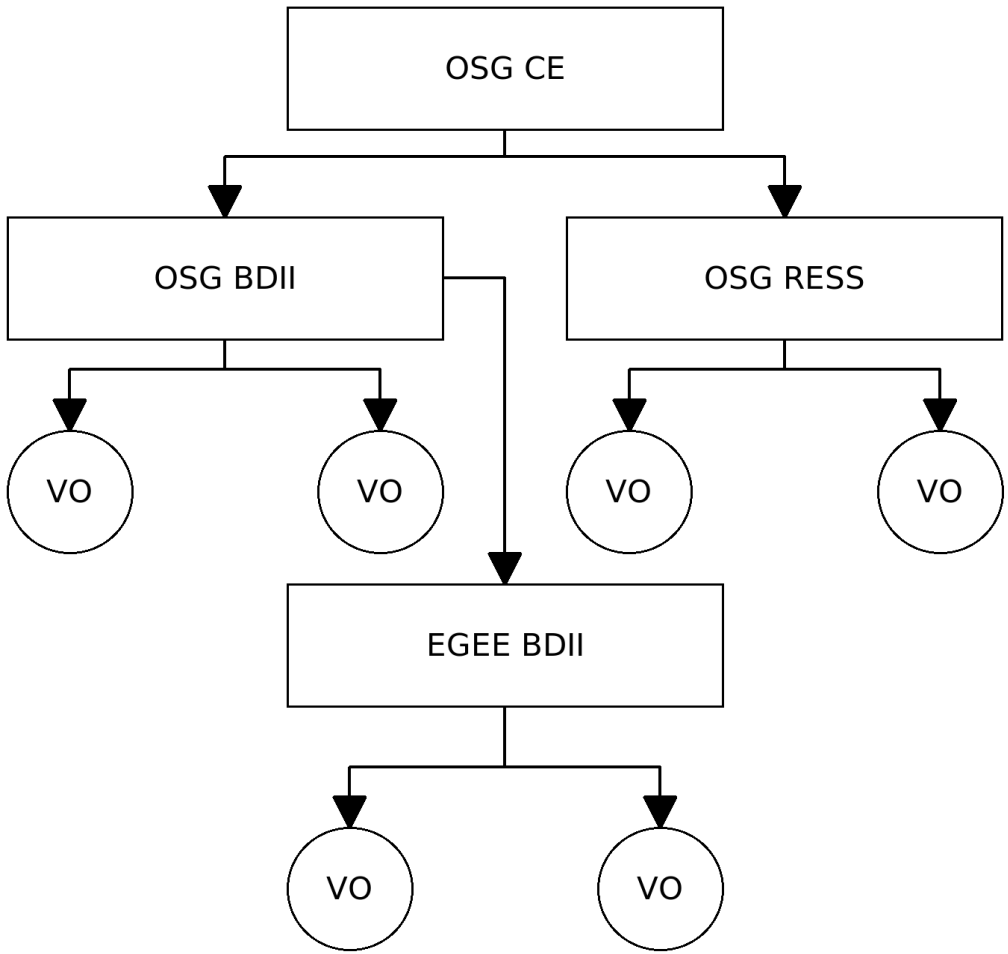

- CE publishes (pushes) information to a BDII Collector
- CE also publishes (pushes) information to ReSS Collector
- VO's access ReSS or the BDII for service/resource discovery

ReSS

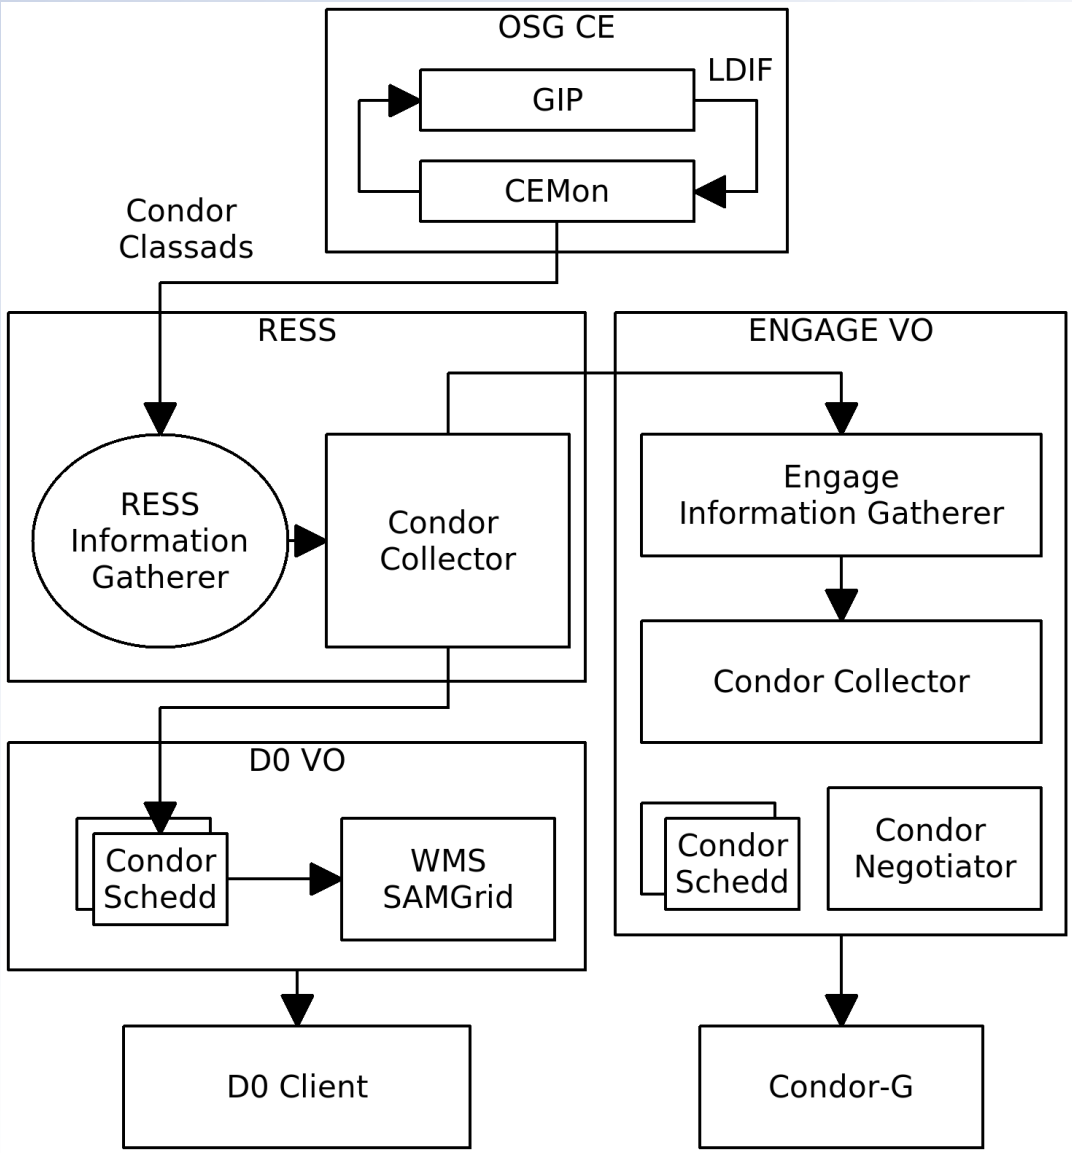

- Pros
  - Simple VO start up
  - No need for a full WMS
  - Submit via Condor-G
- Cons
  - Fragile infrastructure
    - A small (schema legal) change can cause an unmanageable number of extra classads to be generated
    - Single point of failure - service not distributed
    - CEMon can exhibit unstable behavior

BDII

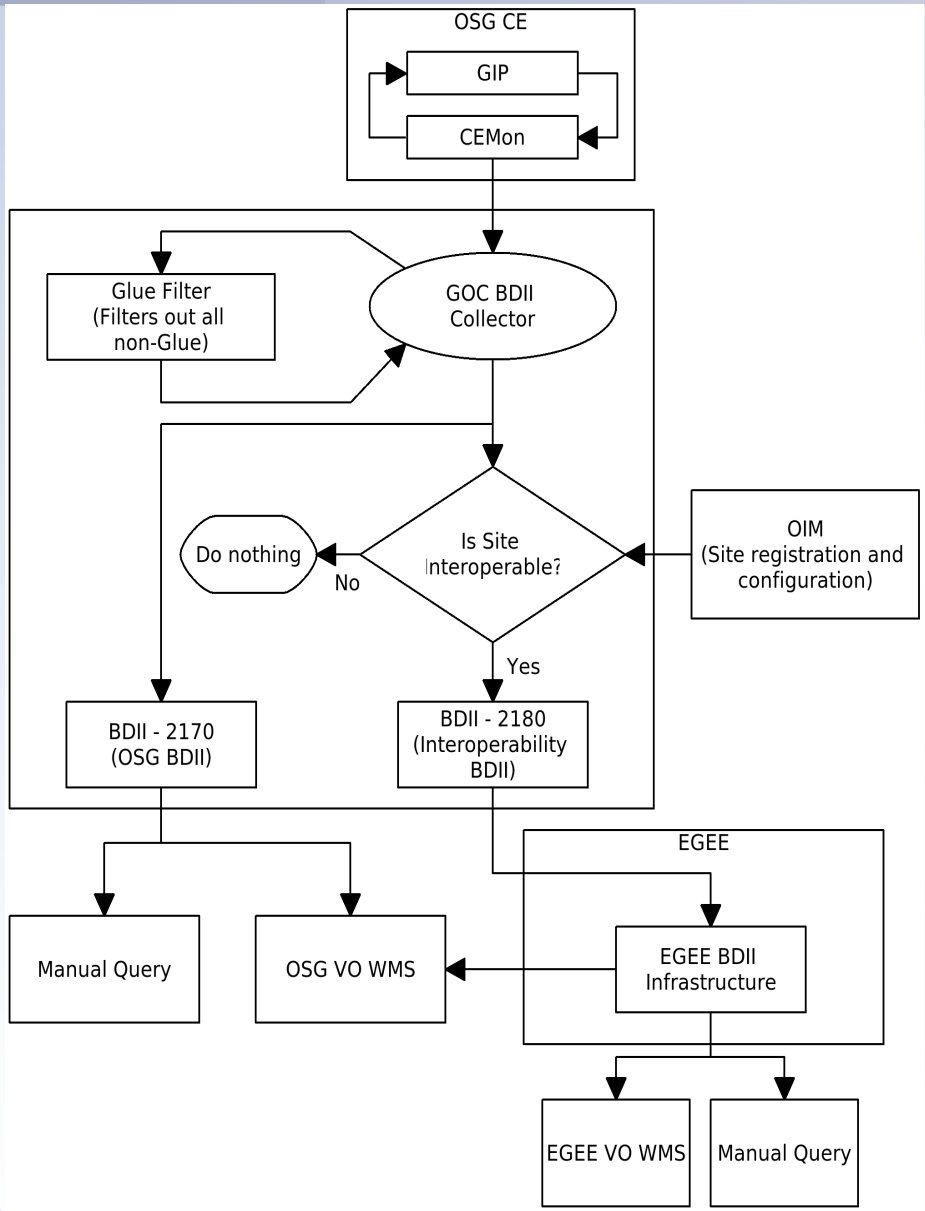

- Pros
  - Interoperability with EGEE
  - Somewhat scalable
  - Very visible information
    - Via ldapsearch
    - Via collector web interface
  - Well defined data structure (Glue Schema)
- Cons
  - Results are machine readable LDIF not very human readable
  - Infrastructure is fragile
    - Adding custom attributes difficult
    - Overloads fairly easily

FTS

- FTS only queries the BDII (WLCG) for Storage Element information
- Information about FTS channels are published to the BDII manually
  - Only one command line tool uses this information
  - According to the developers of FTS, this tool will be deprecated in the next release
  - Only Tier-1's run FTS allowing for quicker upgrade time frames

Case Study - ATLAS

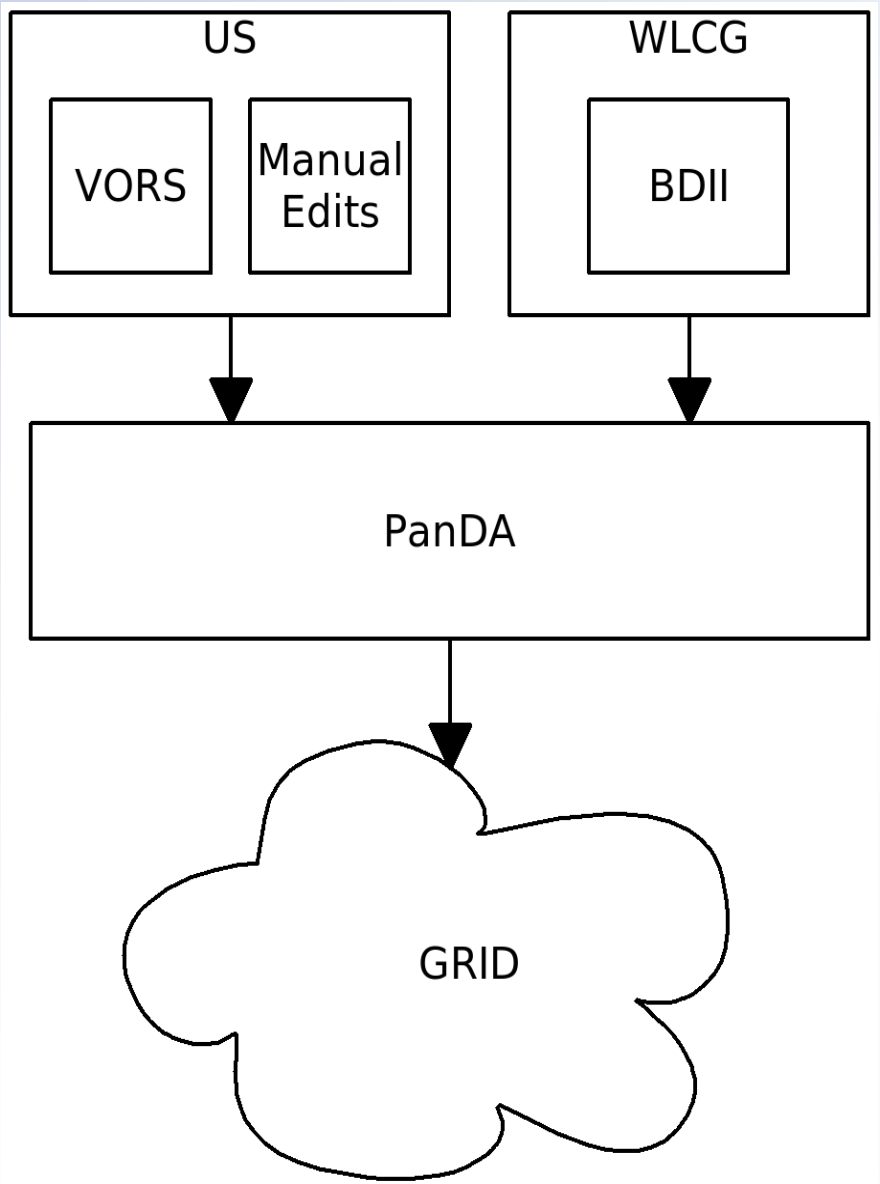

- All information is handled by PanDA
  - PanDA has an internal database of known sites
  - All jobs are matched via the internal database
  - PanDA's database is updated in two ways:
  - EGEE sites - the EGEE BDII is queried
  - US sites - All updates are a manual process
    - A static list of site information is maintained
    - VORS is queried, but does not override static list

Case Study - CMS

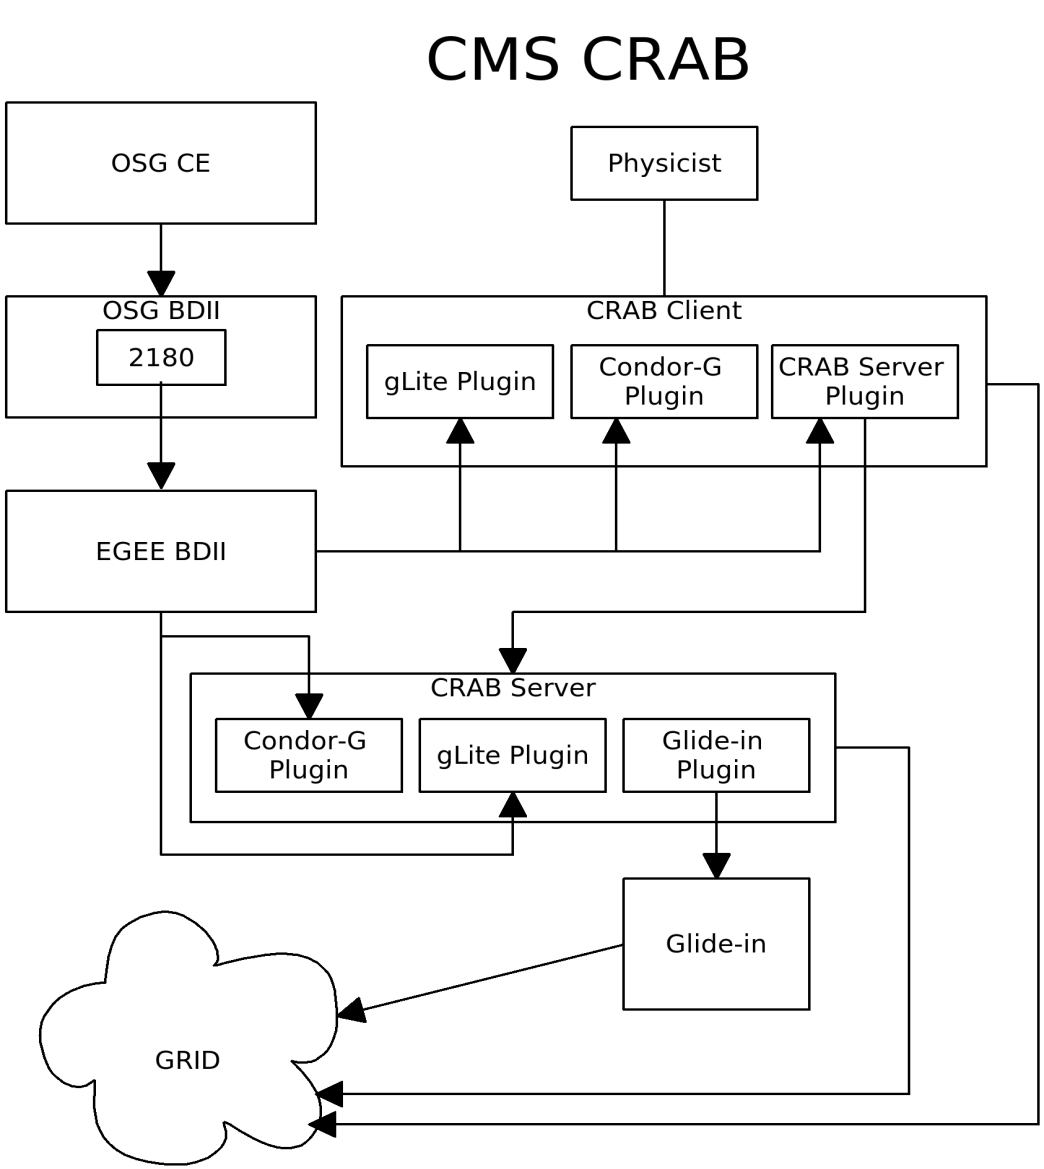

- CMS uses CRAB
  - CRAB client/server only queries the EGEE BDII
  - The EGEE BDII queries the OSG BDII (port 2180) for OSG sites

Critical Functionality

- Must feed LDIF to the WLCG BDII for CMS and ATLAS interoperability
- Must feed to ReSS (Non-WLCG VOs use ReSS)
- A Path is needed to add custom non-Glue attributes to the BDII for various purposes that are un-handled by the current Glue schema.
- Must be Glue 1.3 and Glue 2.0 compatible

Addendum

- OSG maintains an instance of the BDII on port 2170. According to the BDII logs, the only access has been due to monitoring, development and debugging activities.
- The Glide-In WMS only queries the BDII once to get a list of sites to submit jobs to
- ReSS makes assumptions about how schema pieces go together. The assumptions include what is the most important information to generate the combination of classads.
- ReSS has performance limits at 50k classads (no crash).
- ReSS Collector also has web service interface and db backend if configured for access.